



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
27.03.2002 Bulletin 2002/13

(51) Int Cl.7: **G06F 12/06**

(21) Application number: **01308014.8**

(22) Date of filing: **20.09.2001**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
 Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **20.09.2000 US 665989**

(71) Applicant: **Broadcom Corporation**
Irvine, California 92618 (US)

(72) Inventors:
 • **Keller, James B.**
Palo Alto, California 94303 (US)
 • **Hayter, Mark D.**
Menlo Park, California 94025 (US)
 • **Cho, James Y.**
Los Gatos, California 95032 (US)

(74) Representative: **Harris, Ian Richard et al**
D. Young & Co.,
21 New Fetter Lane
London EC4A 1DA (GB)

(54) **Memory controller with programmable configuration**

(57) A memory controller provides programmable flexibility, via one or more configuration registers, for the configuration of the memory. The memory may be optimized for a given application by programming the configuration registers. For example, in one embodiment, the portion of the address of a memory transaction used to select a storage location for access in response to the memory transaction may be programmable. In an implementation designed for DRAM, a first portion may be programmably selected to form the row address and a second portion may be programmably selected to form

the column address. Additional embodiments may further include programmable selection of the portion of the address used to select a bank. Still further, interleave modes among memory sections assigned to different chip selects and among two or more channels to memory may be programmable, in some implementations. Furthermore, the portion of the address used to select between interleaved memory sections or interleaved channels may be programmable. One particular implementation may include all of the above programmable features, which may provide a high degree of flexibility in optimizing the memory system.

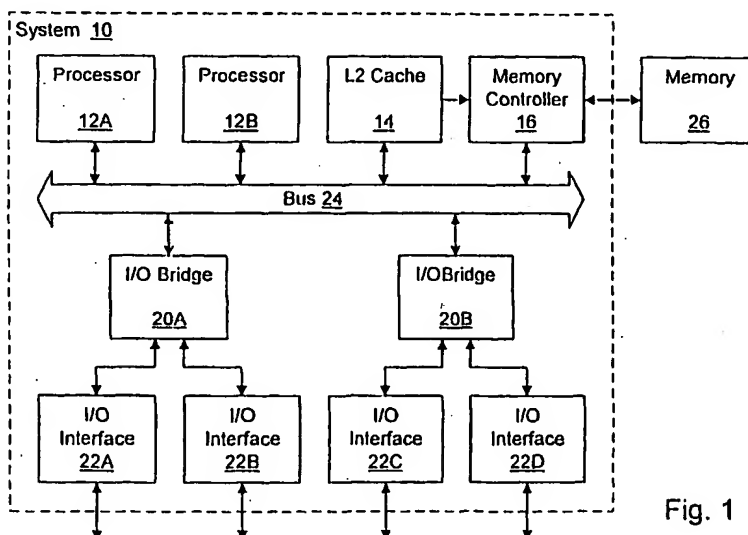


Fig. 1

Description

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] This invention is related to the field of memory controllers.

2. Description of the Related Art

[0002] Memory controllers are generally included in systems to interface various devices in the system with the memory. Generally, the memory controller receives a memory transaction, which includes an address, from one of the devices and accesses the memory using the address. More particularly, the memory controller may use a portion of the address to select a storage location in the memory for access in response to the memory transaction. In a typical dynamic random access memory (DRAM) controller, for example, a first portion of the address is a row address for the DRAM and a second portion of the address is a column address for the DRAM. Together, the row and column addresses select a storage location in the DRAM to be accessed in response to the memory transaction.

[0003] Various memory controllers have implemented other features as well. For example, memory controllers have been configured to attach to multiple memory sections, and the memory controller may provide a separate select signal (typically referred to as a chip select signal) to each memory section. By asserting one of the chip select signals, one of the memory sections may be selected to respond to a memory transaction (e.g. by receiving the row and column addresses provided by the memory controller and reading or writing the selected storage location). Another portion of the address of memory transactions may be used to determine the chip select signals.

[0004] Another feature of some memory controllers is interleaving. With interleaving, two or more memory sections may be combined to represent a given address range. Portions of the data corresponding to the address range may be stored in each of the two or more memory sections, and the memory sections may be accessed in sequence to access all of the data corresponding to the address range. Interleaving may reduce the overall latency of the memory transactions by increasing the possibility of having open pages in the memory (e.g. for a set of transactions clustered in a certain address range) by increasing the number of memory sections accessed at one time. This may effectively increase the page size, which may be viewed as the row size of one memory section multiplied by the number of interleaved memory sections. However, a non-interleaved memory system may provide higher bandwidth, since different transactions can be pipelined into the memory sections (although the latency of each individual transaction may

be higher than an interleaved system). Non-interleaving may allow for more pages (of a smaller size) than interleaved systems. If the different transactions occur to different pages, these transactions may be more likely to be pipelined into the non-interleaved system.

[0005] The various features implemented by memory controllers may attempt to provide good memory performance (e.g. high bandwidth and/or low latency). However, the memory configuration which may deliver the best memory performance may be dependent on the application(s) being executed in a given system. The arrangement of data and instructions used by the application(s) in memory, as well as the pattern of memory transactions performed by the application, may differ from other applications, and the memory configuration which provides the best memory performance may differ from that of other applications. For example, some applications may benefit from an interleaved memory system providing lower latency for each access (e.g. if the memory transactions tend to be clustered in certain address ranges). On the other hand, other applications may benefit from a non-interleaved system (e.g. applications having high numbers of memory transactions, especially if consecutive memory transactions tend to be to disparate addresses). Such applications may benefit from the availability of other memory sections of the non-interleaved configuration to perform the memory transactions. Furthermore, depending upon the arrangement of data in memory, the portions of the address of a memory transaction used to select a storage location in the memory corresponding to the address may affect the performance of the memory system.

[0006] Unfortunately, memory controllers have generally been implemented with a relatively fixed configuration. Typically, address ranges may be assigned to each memory section, and certain memory features (such as page mode) may be enabled or disabled. Otherwise, the configuration of the memory controller (and thus the mapping of addresses to storage locations in the memory) is fixed. Thus, design decisions made when designing the memory controller largely determine the performance that the memory system may deliver for a particular application.

SUMMARY OF THE INVENTION

[0007] Various aspects of the invention are set out in the accompanying claims.

[0008] The problems outlined above are in large part solved by a memory controller as described herein. The memory controller provides programmable flexibility, via one or more configuration registers, for the configuration of the memory. The memory may be optimized for a given application by programming the configuration registers. For example, in one embodiment, the portion of the address of a memory transaction used to select a storage location for access in response to the memory transaction may be programmable. In an implementa-

tion designed for DRAM, a first portion may be programmably selected to form the row address and a second portion may be programmably selected to form the column address. Additional embodiments may further include programmable selection of the portion of the address used to select a bank. By allowing the row address, column address, and bank selection to be programmably selected from the address, the memory system may be optimized by software for the expected memory access patterns. If memory transactions, clustered relatively close in time, are expected to traverse through large blocks of memory, the bank selection may be drawn from more significant bits than the column selection, for example, allowing the memory transactions to experience an open page timing while leaving other banks available for intervening memory transactions to other addresses.

[0009] Still further, interleave modes among memory sections assigned to different chip selects and among two or more channels to memory may be programmable, in some implementations. Furthermore, the portion of the address used to select between interleaved memory sections or interleaved channels may be programmable. Having the interleave, or lack thereof, of memory sections or channels programmable may also provide for flexibility. For example, if accesses to certain address ranges are expected to be accessed with multiple memory transactions clustered close in time, then those address ranges may be represented by two or more interleaved memory sections or channels. Since pages may be open in each of the interleaved memory sections, the interleaving may increase the possibility of receiving transactions which access an open page. On the other hand, other address ranges may not be expected to be accessed with multiple clustered memory transactions, and thus the corresponding memory section or channels may be configured as non-interleaved. The non-interleaved memory sections or channels may have more open pages at different addresses, allowing for different transactions to access an open page.

[0010] One particular implementation may include all of the above programmable features, which may provide a high degree of flexibility in optimizing the memory system.

[0011] Broadly speaking, a memory controller is contemplated for coupling to a memory and for coupling to receive an address of a memory transaction. The memory controller includes one or more registers programmable with an indication of which portion of the address is used to select a storage location in the memory for access in response to the memory transaction. Coupled to the one or more registers and coupled to receive the address, a circuit is configured to extract the portion of the address for transmission to the memory responsive to the indication in the one or more registers.

[0012] Additionally, a system is contemplated. The system comprises a memory and a memory controller coupled thereto. The memory controller is also coupled

to receive an address of a memory transaction. The memory controller is programmable with an indication of which portion of the address is used to select a storage location in the memory for access in response to the transaction, and is configured to extract the portion of the address for transmission to the memory responsive to the indication.

[0013] Moreover, a method is contemplated. A memory controller is programmed with an indication of which portion of an address of a memory transaction is used to select a storage location in a memory for access in response to the memory transaction. The address of the memory transaction is received in the memory controller. The portion of the address is extracted responsive to the indication.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] Embodiments of the invention are described herein, by way of example only, with reference to the accompanying drawings, in which:

Fig. 1 is a block diagram of one embodiment of a system including a memory controller.

Fig. 2 is a block diagram of one embodiment of the memory controller shown in Fig. 1.

Fig. 3 is a block diagram of one embodiment of a channel control circuit shown in Fig. 2 and the channel to a block of memory.

Fig. 4 is a block diagram of one embodiment of the channel control circuit in more detail.

Fig. 5 is a flowchart illustrating one embodiment of chip select generation.

Fig. 6 is a block diagram illustrating an example of various CS interleave modes.

Fig. 7 is a block diagram illustrating an example of channel interleaving.

Fig. 8 is a set of examples according to one embodiment of the memory controller.

Fig. 9 is a block diagram of a carrier medium.

[0015] While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the scope of the claimed invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

System Overview

[0016] Turning now to Fig. 1, a block diagram of one

embodiment of a system 10 is shown. Other embodiments are possible and contemplated. In the embodiment of Fig. 1, system 10 includes processors 12A-12B, an L2 cache 14, a memory controller 16, a pair of input/output (I/O) bridges 20A-20B, and I/O interfaces 22A-22D. System 10 may include a bus 24 for interconnecting the various components of system 10. As illustrated in Fig. 1, each of processors 12A-12B, L2 cache 14, memory controller 16, and I/O bridges 20A-20B are coupled to bus 24. Thus, each of processors 12A-12B, L2 cache 14, memory controller 16, and I/O bridges 20A-20B may be an agent on bus 24 for the illustrated embodiment. I/O bridge 20A is coupled to I/O interfaces 22A-22B, and I/O bridge 20B is coupled to I/O interfaces 22C-22D. L2 cache 14 is coupled to memory controller 16, which is further coupled to a memory 26.

[0017] Processors 12A-12B may be designed to any instruction set architecture, and may execute programs written to that instruction set architecture. Exemplary instruction set architectures may include the MIPS instruction set architecture (including the MIPS-3D and MIPS MDMX application specific extensions), the IA-32 or IA-64 instruction set architectures developed by Intel Corp., the PowerPC instruction set architecture, the Alpha instruction set architecture, the ARM instruction set architecture, or any other instruction set architecture. While system 10 as shown in Fig. 1 includes two processors, other embodiments may include one processor or more than two processors, as desired.

[0018] L2 cache 14 is a high speed cache memory. L2 cache 14 is referred to as "L2" since processors 12A-12B may employ internal level 1 ("L1") caches. If L1 caches are not included in processors 12A-12B, L2 cache 14 may be an L1 cache. Furthermore, if multiple levels of caching are included in processors 12A-12B, L2 cache 14 may be an outer level cache than L2. L2 cache 14 may employ any organization, including direct mapped, set associative, and fully associative organizations. In one particular implementation, L2 cache 14 may be a 512 kilobyte, 4 way set associative cache having 32 byte cache lines. A set associative cache is a cache arranged into multiple sets, each set comprising two or more entries. A portion of the address (the "index") is used to select one of the sets (i.e. each encoding of the index selects a different set). The entries in the selected set are eligible to store the cache line accessed by the address. Each of the entries within the set is referred to as a "way" of the set. The portion of the address remaining after removing the index (and the offset within the cache line) is referred to as the "tag", and is stored in each entry to identify the cache line in that entry. The stored tags are compared to the corresponding tag portion of the address of a memory transaction to determine if the memory transaction hits or misses in the cache, and is used to select the way in which the hit is detected (if a hit is detected).

[0019] Memory controller 16 is configured to access memory 26 in response to memory transactions re-

ceived on bus 24. Memory controller 16 receives a hit signal from L2 cache 14, and if a hit is detected in L2 cache 14 for a memory transaction, memory controller 16 does not respond to that memory transaction. Other embodiments may not include L2 cache 14 and memory controller 16 may respond to each memory transaction. If a miss is detected by L2 cache 14, or the memory transaction is non-cacheable, memory controller 16 may access memory 26 to perform the read or write operation. Memory controller 16 may be designed to access any of a variety of types of memory. For example, memory controller 16 may be designed for synchronous dynamic random access memory (SDRAM), and more particularly double data rate (DDR) SDRAM. Alternatively, memory controller 16 may be designed for DRAM, Rambus DRAM (RDRAM), SRAM, or any other suitable memory device.

[0020] I/O bridges 20A-20B link one or more I/O interfaces (e.g. I/O interfaces 22A-22B for I/O bridge 20A and I/O interfaces 22C-22D for I/O bridge 20B) to bus 24. I/O bridges 20A-20B may serve to reduce the electrical loading on bus 24 if more than one I/O interface 22A-22B is bridged by that I/O bridge. Generally, I/O bridge 20A performs transactions on bus 24 on behalf of I/O interfaces 22A-22B and relays transactions targeted at an I/O interface 22A-22B from bus 24 to that I/O interface 22A-22B. Similarly, I/O bridge 20B generally performs transactions on bus 24 on behalf of I/O interfaces 22C-22D and relays transactions targeted at an I/O interface 22C-22D from bus 24 to that I/O interface 22C-22D. In one implementation, I/O bridge 20A may be a bridge to a PCI interface (e.g. I/O interface 22A) and to a Lightning Data Transport (LDT) I/O fabric developed by Advanced Micro Devices, Inc. (e.g. I/O interface 22B). Other I/O interfaces may be bridged by I/O bridge 20B. Other implementations may bridge any combination of I/O interfaces using any combination of I/O bridges. I/O interfaces 22A-22D may include one or more serial interfaces, Personal Computer Memory Card International Association (PCMCIA) interfaces, Ethernet interfaces (e.g. media access control level interfaces), Peripheral Component Interconnect (PCI) interfaces, LDT interfaces, etc.

[0021] Bus 24 may be a split transaction bus, in one embodiment. Bus 24 may employ a distributed arbitration scheme, in one embodiment. In one embodiment, bus 24 may be pipelined. Bus 24 may employ any suitable signalling technique. For example, in one embodiment, differential signalling may be used for high speed signal transmission. Other embodiments may employ any other signalling technique (e.g. TTL, CMOS, GTL, HSTL, etc.).

[0022] It is noted that system 10 (and more particularly processors 12A-12B, L2 cache 14, memory controller 16, I/O interfaces 22A-22D, I/O bridges 20A-20B and bus 24 may be integrated onto a single integrated circuit as a system on a chip configuration. In another configuration, memory 26 may be integrated as well. Alterna-

tively, one or more of the components may be implemented as separate integrated circuits, or all components may be separate integrated circuits, as desired. Any level of integration may be used.

[0023] It is noted that, while the illustrated embodiment employs a split transaction bus with separate arbitration for the address and data buses, other embodiments may employ non-split transaction buses arbitrated with a single arbitration for address and data and/or a split transaction bus in which the data bus is not explicitly arbitrated. Either a central arbitration scheme or a distributed arbitration scheme may be used, according to design choice. Furthermore, bus 24 may not be pipelined, if desired.

[0024] It is noted that, while Fig. 1 illustrates I/O interfaces 22A-22D coupled through I/O bridges 20A-20B to bus 24, other embodiments may include one or more I/O interfaces directly coupled to bus 24, if desired.

[0025] As used herein, the term "memory transaction" refers to a communication between a device and the memory system. The memory transaction includes an address identifying a storage location in the memory. In the illustrated embodiment, the memory transaction is transmitted on bus 24 to memory controller 16. However, other systems may use other communications media (e.g. packet based transmission, clock-forwarded links, point to point interconnect, etc.). Memory transactions may include read transactions (a transfer of data from memory to the device) and write transactions (a transfer of data from the device to memory).

Memory Controller

[0026] Turning now to Fig. 2, a block diagram of one embodiment of memory controller 16 is shown. Other embodiments are possible and contemplated. In the embodiment of Fig. 2, memory controller 16 includes a transaction queue 30 and a pair of channel control circuits 32A-32B. Transaction queue 30 is coupled to receive memory transactions from bus 24, and is coupled to channel control circuits 32A-32B. Channel control circuit 32A is coupled to a first channel 34A and channel control circuit 32B is coupled to a second channel 34B. Channel control circuit 32A includes a set of one or more configuration registers 36A, and channel control circuit 32B includes a set of one or more configuration registers 36B.

[0027] Generally, channel control circuits 32A-32B access memory attached to the corresponding channel 34A-34B in response to memory transactions provided from transaction queue 30. Configuration registers 36A-36B are used to program the configuration of the memory system. Configuration registers 36A-36B may be made visible to software so that the configuration may be programmed. For example, configuration registers 36A-36B may be mapped to certain addresses within the addressable range provided by bus 24 ("memory mapped"), and thus may be programmed by executing

store instructions in processor 12A or 12B to the addresses mapped to the registers. Configuration registers 36A-36B may allow for a high degree of flexibility in configuring the memory, thus allowing optimization of the memory system for the applications to be executed on system 10. Additional details regarding the configuration options of one embodiment of memory controller 16 are provided in more detail below.

[0028] As illustrated in Fig. 2, memory controller 16 may employ multiple channels. Other embodiments are contemplated employing more channels, as well as embodiments having only a single channel. Each channel control circuit 32A-32B performs memory transactions mapped to the corresponding channel. Additionally, each channel control circuit 32A-32B includes a set of configuration registers 36A-36B. Accordingly, the memory attached to each channel may be configured differently. Additional flexibility in the memory system configuration may be provided via the use of different configuration registers for each channel.

[0029] In one embodiment, the two channels may be configurable as interleaved or not interleaved via configuration registers 36A-36B. Non-interleaved channels may provide high bandwidth, as the two channels may be independent of each other and thus may respond to memory transactions in parallel. Interleaved channels may increase the possibility of finding an open page if memory transactions are clustered in certain address ranges, thus decreasing the latency of the memory transactions. Embodiments employing more than two channels may interleave two or more of the channels while having remaining channels non-interleaved, via configuration registers 36A-36B and other configuration registers in the additional channel control circuits.

[0030] Transaction queue 30 is configured to receive and queue memory transactions from bus 24, and to issue those transactions to one of channel control circuits 32A-32B. Transaction queue 30 may have access to at least the configuration registers 36A-36B that define which addresses are mapped to which memory channel, to select the channel control circuit to which a given transaction is issued. Alternatively, transaction queue 30 may include one or more shadow registers storing indications which represent the addresses that are mapped to each memory channel, or may include a separate set of programmable registers. Transaction queue 30 may also receive and queue the data for write transactions, and provide that data to the corresponding channel control circuit 32A-32B for updating the memory. Transaction queue 30 may receive and queue data from the corresponding channel control circuit 32A-32B for read transactions, and provide the data on bus 24 for consumption by the device which initiated the memory transaction.

[0031] As used herein, the term "channel" refers to a communication path to a memory. The communication path may include the interface signals used to interface to that memory, and is independent of other channels.

[0032] As used herein, the term "registers" refer to any clocked storage device. For example, registers may include clocked RAM cells, latches, flip flops, etc.

[0033] Turning next to Fig. 3, a block diagram of one embodiment of channel control circuit 32A and channel 34A is shown. Other embodiments are possible and contemplated. Channel 34B may be similar, in one embodiment. The embodiment of Fig. 3 may be suitable for interfacing to SDRAM memory, such as one or more dual inline memory modules (DIMMs) 40A-40D. Each of DIMMs 40A-40D may be populated by one or more SDRAM chips (e.g. SDRAM chips 42A-42D illustrated on DIMM 40A). DIMMs 40A-40D (and optionally one or more DIMMs on channel 34B) may be one embodiment of memory 26 shown in Fig. 1.

[0034] Channel 34A as shown in Fig. 3 includes signals for interfacing to DIMMs 40A-40D. More particularly, channel 34A includes a separate chip select line carrying a chip select signal (CS0-CS3) for each DIMM 40A-40D. The chip select signal may be asserted to select the corresponding DIMM 40A-40D to respond to a memory transaction. Mappings of address ranges to chip select signals may be programmed into configuration registers 36A. While each chip select signal is connected to one DIMM 40A-40D, the remaining signals of memory channel 34A may be connected to each of DIMMs 40A-40D.

[0035] Channel 34A includes a clock line or lines (CLK) for transmitting a clock signal to DIMMs 40A-40D. Memory channel 34A also includes data transmission lines including lines for the data (e.g. 64 bits of data, in one embodiment), optional error checking and correction code (ECC) information corresponding to the data, and data strobe signals corresponding to each byte of data and the ECC information. Row address strobe (RAS), column address strobe (CAS) and write enable signals are included as well. These signals are used to transmit commands to the DIMM 40A-40D selected using the chip select signals. More particularly, the commands may include an activate command to activate a row of storage locations in a bank. The activate command may include the row address on the address bus (A[12:0]). One or more read or write commands may be issued to the active bank, and each read or write command may include the column address being transmitted on the address bus. Other commands may include a precharge command to close the active bank, a burst terminate command to terminate a burst of data starting at the identified column, etc. Finally, channel 34A includes a bank select signals (BA[1:0]) used to select a bank within the SDRAMs within the selected DIMM 40A-40D. The bank may be transmitted with both the row address and the column address.

[0036] While Fig. 3 illustrates one or more DIMMs forming the memory of a channel, other embodiments may form a channel from single inline memory modules (SIMMs), individual RAM chips, etc. Additionally, a single DIMM may receive more than one chip select signal,

or multiple DIMMs, SIMMs, or individual RAM chips may share a chip select (and be coupled to different portions of the data bus).

[0037] As used herein, the term "row" refers to a set of storage locations in a DRAM, DRAMs are configured with multiple rows, and one row is selected for a transaction by transmitting a row address to the DRAM (e.g. with an activate command on an SDRAM or an assertion of the RAS signal for asynchronous DRAM). The term "column" refers to a storage location within the set of storage locations in each row. One column is selected for a transaction by transmitting a column address to the DRAM (e.g. with a read or write command on an SDRAM or the assertion of the CAS signal for asynchronous DRAM). Thus, DRAM memory is arranged into rows and columns of storage locations forming a two dimensional array of storage locations. A given storage location is selected using a particular row address and a particular column address. The intersection of that row and column within the two dimensional array is the storage location. The number of bits or bytes in the selected storage location may vary from embodiment to embodiment. Furthermore, multiple storage locations may be accessed in parallel to provide one transfer on the data lines between the memory and memory controller 16, and multiple transfers on the data lines (each affecting one or more storage locations in the DRAM memory) may be performed to access all of the data corresponding to a memory transaction.

[0038] Multiple reads and/or writes to various columns in the active row may be performed without retransmitting the row address. Reads or writes which are accomplished without retransmitting the row address are often referred to as page mode accesses or referred to as hitting an open page.

[0039] A given SDRAM may include multiple sets of rows and columns of storage locations. The term "bank" refers to one set of rows and columns. The bank is selected by transmitting the bank signals during the activate and read/write commands. Other memories besides SDRAMs may include banks of rows and columns as well.

[0040] Turning next to Fig. 4, a block diagram of one embodiment of channel control circuit 32A is shown. Other embodiments are possible and contemplated. Channel control circuit 32B may be similar. In the embodiment shown in Fig. 4, channel control circuit 32A includes a chip select (CS) generator circuit 50, a row/column/bank extract circuit 52, and a memory interface control circuit 54. CS generator circuit 50 and row/column/bank extract circuit 52 are coupled to receive an address of a memory transaction being issued to channel control circuit 32A from transaction queue 30 and are further coupled to memory interface control circuit 54 and to each other. Memory interface control circuit 54 is further coupled to channel 34A and to transaction queue 30. As illustrated in Fig. 4, CS generator circuit 50 includes configuration registers 36AA, 36AB, 36AC,

and 36AD. Extract circuit 52 includes configuration registers 36AE, 36AF, 36AG, 36AH, 36AI, 36AJ, 36AK, 36AL, 36AM, 36AN, 36AO, and 36AP. Memory interface control circuit 54 includes configuration register 36AQ. Configuration registers 36AA-36AQ may be included in one embodiment of configuration registers 36A.

[0041] Extract circuit 52 is configured to extract portions of the address of a memory transaction for locating a storage location in the memory attached to channel 34A to be accessed in response to the memory transaction. Extract circuit 52 provides the extracted portions to memory interface control circuit 54 for use in communicating on channel 34A. For example, the illustrated embodiment may be used for SDRAM memory. Thus, portions of the address may be extracted to provide the row address, column address, and bank selection to the memory.

[0042] Via configuration registers 36AE-36AP, extract circuit 52 may be programmable to select which portions of the address are used for the row address, column address, and bank select. For example, in the illustrated embodiment, the portions of the address which are extracted as the row address, column address, and bank selection are separately programmable for each chip select signal supported by channel control circuit 52. Thus, configuration register 36AE is programmed with an indication of which portion of the address is extracted for the row address if CS0 is the chip select signal corresponding to the address. Similarly, configuration registers 36AF-36AH are programmed with indications of which portion of the address is extracted for the row address if CS1, CS2, or CS3 is the chip select signal corresponding to the address, respectively. Configuration register 36AI is programmed with an indication of which portion of the address is extracted for the column address if CS0 is the chip select signal corresponding to the address. Similarly, configuration registers 36AJ-36AL are programmed with indications of which portion of the address is extracted for the row address if CS1, CS2, or CS3 is the chip select for the address, respectively. Finally, configuration register 36AM is programmed with an indication of which portion of the address is extracted for the bank select if CS0 is the chip select signal corresponding to the address. Similarly, configuration registers 36AN-36AP are programmed with indications of which portion of the address is extracted for the bank select if CS1, CS2, or CS3 is the chip select signal corresponding to the address, respectively.

[0043] Extract circuit 52 may receive an indication of which chip select signal corresponds to the address of the memory transaction from CS generator circuit 50, and may extract the row address, column address, and bank selection from the address using the corresponding indications from configuration registers 36AE-36AP. For example, in one embodiment, each of configuration registers 36AE-36AP may store a bit mask having a bit assigned to each address bit. If the bit in the bit mask is

set, the corresponding address bit is extracted and used in the row address, column address, or bank selection. If the bit in the mask is clear, the corresponding address bit is not extracted and used in the row address, column address, or bank selection. Alternatively, the bit in the bit mask being clear may indicate extraction of the corresponding address bit and the bit in the bit mask being set may indicate lack of extraction of the corresponding address bit.

[0044] In one particular implementation, each of the bit masks are required to identify contiguous address bits for inclusion in the corresponding row address, column address, or bank selection. The address bits used to form the row address need not be contiguous to the address bits used to form the column address or the bank selection, and the address bits used to form the column address need not be contiguous to the address bits used to form the row address or the bank selection. In such an implementation, extract circuit 52 may apply the bit mask to the address to mask off the non-selected address bits. Additionally, extract circuit 52 may right shift the masked address such that the least significant of the selected address bits is the least significant bit of the resulting value. The shifting and masking may be performed serially (in either order) or in parallel.

[0045] As an alternative to bit masks, configuration registers 36AE-36AP may be programmed in other fashions to indicate the portions of the address to be used as the row address, column address, and bank selection. For example, embodiments which require that contiguous address bits form a given row address, column address, or bank selection may be programmed with numbers indicating the least significant bit and most significant bit of the row address, column address, or bank selection. In yet another alternative, a number indicative of either the least significant bit or most significant bit and a bit mask selecting the number of address bits used by the corresponding memory may be programmed into configuration registers 36AE-36AP, as desired.

[0046] In another implementation, each of the bit masks are required to identify contiguous address bits for inclusion in the corresponding row address or bank selection. The address bits used to form the row address need not be contiguous to the address bits used to form the column address or the bank selection, and the address bits used to form the column address need not be contiguous to the address bits used to form the row address or the bank selection. As to the column address, the bit mask may be required to have two contiguous sections, with a gap of one or more bits between the sections. In one specific implementation, one of the sections may comprise bit 5 or bits 5 and 6 of the address (where the least significant address bit is referred to as bit 0 of the address) and the other section may comprise more significant bits. Such an implementation may allow for memory transactions to consecutive cache lines (for which address bits may vary in bit 5, or bits 5 and 6, only

and thus may have the same row address and bank selection, but different column addresses) to be more likely to hit an open page in the memory, thus reducing the overall latency of the memory transactions. Such memory transactions may occur if a device frequently access blocks of memory which are larger than a cache line. The address bits between the address bits in the first section and the second section of the column address may be used for bank selection, for interleave selection among the memory sections assigned to different chip select signals (described in more detail below), for interleave selection between the channels (described in more detail below), etc.

[0047] CS generator circuit 50 is configured to generate the chip select signals for each memory transaction. The generated chip select signals are provided to memory interface control circuit 54 for transmission on channel 34A and to row/column/bank extract circuit 52. More particularly, CS generator circuit 50 is programmable, via configuration registers 36AA-36AD, to generate chip select signals. Configuration register 36AC is programmed with an indication of the starting address of a range of addresses mapped to each chip select, using respective fields CS0_Start, CS1_Start, CS2_Start, and CS3_Start as illustrated in Fig. 4. More particularly, in one embodiment, the most significant bits of the starting address may be programmed into the corresponding start field. The number of most significant bits provided in each field may be determined based on the minimum memory size which is supported for each chip select signal. For example, if the minimum supported memory size on a given chip select signal were 16 Megabytes, address bits more significant than bit 23 (where bit 0 is the least significant address bit) would be included in each field. Similarly, configuration register 36AD is programmed with an indication of the ending address of the range of addresses using respective fields CS0_End, CS1_End, CS2_End, and CS3_End as illustrated in Fig. 4. More particularly, in one embodiment, the most significant bits of the ending address plus 1 may be programmed into the corresponding end field. Thus, a given chip select signal may be decoded (subject to certain options described below) for an address if the most significant bits of the address are greater than or equal to the most significant bits of the start address for that given chip select signal and less than the most significant bits of the end address for that given chip select signal.

[0048] Additional flexibility in memory configuration may be provided via configuration registers 36AA and 36AB. Configuration register 36AB includes a CS mode field which may be programmed with various mode options for the chip select signals managed by channel control circuit 32A. For example, in one embodiment, portions of memory corresponding to two or more chip select signals may be interleaved while other portions are not interleaved. To interleave the memory corresponding to two chip select signals, the corresponding start and end addresses in configuration registers

36AC-36AD may be programmed to the same values (e.g. an address range large enough to cover the amount of memory assigned to both chip select signals), and an indication that the two chip select signals are interleaved may be programmed into the CS mode field of configuration register 36AB. In one implementation, channel control circuit 32A may provide for any of the following interleave options via encodings in the CS mode field: (i) no interleave; (ii) CS0 and CS1 memory sections interleaved, CS2 and CS3 memory sections not interleaved; (iii) CS1 and CS2 memory sections interleaved, CS0 and CS3 memory sections not interleaved; (iv) CS2 and CS3 memory sections interleaved, CS0 and CS1 memory sections not interleaved; (v) CS0, CS1, CS2, and CS3 memory sections interleaved. For example, in one embodiment, the CS mode field may comprise a four bit value, each bit assigned to one of the chip select signals. If the assigned bit is set, the memory section corresponding to that chip select signal is interleaved with the memory sections corresponding to other chip select signals for which the assigned bit in the CS mode field is also set. If the assigned bit is clear, the memory section corresponding to that chip select signal is not interleaved.

[0049] If the memory sections corresponding to two or more chip select signals are interleaved, configuration register 36AA may be programmed with an indication of which portion of the address is used to select between the interleaved memory sections. For example, configuration register 36AA may store a bit mask having set bits (or alternatively clear bits) for the bit or bits which identify which interleaved memory section is selected in response to an address, similar to the row, column, and bank select bit masks. Alternatively, bit numbers may be programmed into configuration register 36AA. For example, if memory sections corresponding to two chip select signals are interleaved, one bit of the address may be used to select the corresponding memory section and thus the chip select signal to activate. If memory sections corresponding to four chip select signals are interleaved, two bits of the address may be used to select the corresponding memory section and thus the chip select signal to activate.

[0050] In one embodiment, if the memory sections corresponding to two or more chip selects are interleaved, the other configuration parameters may be set the same (e.g. row, column, and bank configurations in configuration registers 36AE-36AP may be set the same).

[0051] Channel control circuit 32A and more particularly CS generator circuit 50 may support channel interleaving as well. The Channel_Sel field of configuration register 36AB may store an indication of whether or not the channels are interleaved and an indication of the portion of the address used to select between the interleaved channels. In one embodiment in which memory controller 16 includes two channels, the Channel_Sel field may be a number which, if zero, indicates no chan-

nel interleave and, if non zero, indicates channel interleave and also is the bit number of the address which selects between the channels. Other embodiments may have more than two channels which may be interleaved, and the Channel_Sel field may be similar to the row indication in configuration register 36AE. Still further, the Channel_Sel field may comprise an indication of channel interleaving and a separate value (e.g. bit mask or bit numbers). If the channels are interleaved, other memory parameters may be set the same between the two channels (e.g. the remainder of configuration registers 36AA-36AQ and corresponding registers in configuration registers 36B may be set the same). Fig. 5 below is a flowchart illustrating chip select generation in more detail.

[0052] Memory interface control circuit 54 generally is configured to sequence commands on channel 34A to perform the desired memory transaction. Thus, memory interface control circuit 54 may receive additional transaction information from transaction queue 30, such as the size of the transaction, the read/write nature of the transaction, etc. Memory interface control circuit 54 may comprise a queue of in-progress transactions, with a scoreboard of commands to be sequenced at various times (e.g. measured on a clock cycle basis). The scoreboard may be shifted each clock cycle, and commands performed based on the state of the scoreboard.

[0053] Memory interface control circuit 54 may determine whether or not to close an active page after a transaction is complete. Memory interface control circuit 54 may be programmable, via configuration register 36AQ, with a page mode policy for the memory section corresponding each chip select signal. In one embodiment, the available page mode policies may include a close page policy in which the page is closed at the end of each transaction or a CAS time check policy in which transaction queue 30 is checked as the read or write command for a transaction is being issued on memory channel 34A. If other transactions to the same page are in transaction queue 30, the page may be kept open. Otherwise, the page may be closed. In yet another embodiment, the page mode policy may be based on a hint in the transaction. The page mode policy may be programmed differently for the memory corresponding to each of the chip selects, and thus configuration register 36AQ may include page mode policy fields CS0_PM, CS1_PM, CS2_PM, and CS3_PM for chip selects CS0-CS3, respectively.

[0054] As the above discussion illustrates, memory controller 16 may provide for a large amount of flexibility in configuring the memory to which memory controller 16 is coupled. For example, by allowing the row address, column address, and bank selection to be programmably selected from the address, the memory system may be optimized by software for the expected memory access patterns. If memory transactions, clustered relatively close in time, are expected to traverse through large blocks of memory, the bank selection may

be drawn from more significant bits than the column selection, for example, allowing the memory transactions to experience an open page timing while leaving other banks available for intervening memory transactions to other addresses.

[0055] Having the interleave, or lack thereof, of memory sections corresponding to different chip select signals or channels programmable also provides for flexibility. For example, if accesses to certain address ranges are expected to be accessed with multiple memory transactions clustered close in time, then those address ranges may be represented by two or more interleaved memory sections corresponding to two or more interleaved chip select signals. Since pages may be open in the memory corresponding to each of the interleaved chip select signals, the interleaving may increase the possibility of receiving transactions which access an open page. On the other hand, other address ranges may not be expected to be accessed with multiple clustered memory transactions, and thus the corresponding memory may be configured as non-interleaved. The non-interleaved memory sections may have more open pages at different addresses, allowing for different transactions to access an open page. Similarly, interleaving channels may further increase the number of open pages in the affected memory range. On the other hand, higher memory bandwidth may be achieved by not interleaving the channels. Accordingly, software may optimize the configuration of memory based on the expected memory access pattern.

[0056] It is noted that, while the above discussion of Fig. 4 has described an embodiment for interfacing to SDRAM, other embodiments may interface to asynchronous DRAM, non-banked DRAM (including only row address and column address but no bank selection), SRAM, etc. Any suitable type of memory may be interfaced, and the configurability of memory controller 16 may be modified accordingly. Generally, some embodiments of memory controller 16 may be programmed to select at least one portion of the address for selecting a memory location in the memory using configuration registers 36A-36B.

[0057] It is noted that, while the illustrated embodiment provides for programmability of row address, column address, bank selection, and page mode policy on a chip select granularity, other embodiments may provide for programmability of one or more of these attributes at the channel granularity instead, in other embodiments.

[0058] It is noted that, while one implementation above includes a requirement that the various bit masks (e.g. row address bit mask, column address bit mask, etc.) select contiguous address bits for inclusion in the corresponding value, other embodiments may allow any random set of address bits to be selected.

[0059] It is noted that, while Fig. 4 illustrates a specific allocation of values to certain configuration registers 36AA-36AQ, the information stored may be rearranged

into any combination of registers according to design choice. It is further noted that, while the illustrated embodiment supports 4 chip select signals, other embodiments may support any number of chip select signals, according to design choice.

[0060] It is still further noted that, while the embodiment of Fig. 4 provides for programmability of a variety of memory system features, other embodiments may employ any subset of the programmability, as desired. For example, chip select interleave could be programmable while row address, column address, and bank selection could be fixed based on memory size. Alternatively, chip select interleave might not be programmable (and might be interleaved, partially interleaved, or not interleaved), but row address, column address, and bank selection may be programmable.

[0061] It is noted that, while CS generator circuit 50, extract circuit 52, and memory interface control circuit 54 are shown as separate circuits in Fig. 4 for ease of understanding the programmability of memory controller 16, these circuits may be integrated together in implementation, if desired.

[0062] As used herein, the term "bit mask" refers to a binary value in which each bit corresponds to one or more bits of a first value to be masked by the bit mask. The bit being in one state (set or clear) may indicate a masking off of the one or more bits of the first value in the masking operation, while the bit being in the other state may indicate passing through of the one or more bits of the first value in the masking operation.

[0063] As used herein, the term "memory section" refers to a one or more memory circuits which are treated as a unit by memory controller 16. Thus, the memory circuits which receive a particular chip select signal may be a memory section, in one example. The memory circuits on one channel of memory controller 16 may be a memory section, in another example. Generally, a memory circuit may be any type of memory circuit, including DIMMs, SIMMs, memory chips directly mounted on a circuit board to which memory controller 16 (or system 10) is mounted, etc. The memory circuits may be SDRAM, DRAM, or any other form of memory, as desired.

[0064] As used herein, two or more memory sections are "interleaved" if an address region is mapped to the combination of the memory sections and the address range is divided into blocks, where the number of blocks is at least twice the number of memory sections. The blocks are alternately mapped to each of the interleaved memory sections. The number of blocks may be dependent on which address bit or bits is used to select between the interleaved portions as compared to the bits used to form the row address, column address, etc.

[0065] As used herein, the term "page mode policy" refers to the policy used to determine whether or not to keep a page accessed by a memory transaction open after the transaction is performed. Generally, a page refers to the amount of memory made accessible by pre-

senting one row address to the memory (e.g. via an activate command in an SDRAM or via assertion of the RAS signal in an asynchronous DRAM). A page is open if the row address has been transmitted and the page has not been closed (e.g. with a precharge command or via autoprecharge in an SDRAM or by deassertion of the RAS signal in an asynchronous DRAM). Accesses to an open page present the column address to the memory, but need not perform the row address transfer, and thus may be performed with lower latency than if the row address were to be presented as well.

[0066] Turning next to Fig. 5, a flowchart is shown illustrating generation of chip select signals for channel control circuit 32A according to one embodiment of memory controller 16. Other embodiments are possible and contemplated. While the blocks shown in Fig. 5 are illustrated in a particular order for ease of understanding, any suitable order may be used. Furthermore, blocks may be performed in parallel in combinatorial logic circuitry within memory controller 16.

[0067] Memory controller 16 determines if channel interleaving is programmed into configuration registers 36A-36B (decision block 60). More particularly, memory controller 16 may examine the Channel_Sel field of configuration register 36AB. If channel interleaving is selected, memory controller 16 further determines if the address bit used to select between the interleaved channels is clear (decision block 62). If the address bit is not clear, the memory transaction is mapped to channel control circuit 32B and thus no chip selects in channel control circuit 32A are generated.

[0068] On the other hand, if channel interleaving is not programmed into configuration registers 36A-36B or the address bit used to select between the interleaved channels is clear, then channel control circuit 32A may handle the memory transaction. The address of the memory transaction is compared to the CS address ranges programmed into configuration registers 36AC-36AD (block 64). Memory controller 16 determines if one or more of the CS ranges include the address of the memory transaction (decision block 66). If none of the CS ranges encompasses the address, then channel control circuit 32A does not handle the memory transaction. On the other hand, if at least one CS range encompasses the address, then channel control circuit 32A does handle the memory transaction.

[0069] Memory controller 16 determines if the CS mode is interleaved for the CS range(s) encompassing the address of the memory transaction (decision block 68). If the CS mode is interleaved, then the CS interleave indication in configuration registers 36AA specifies which chip select signal to activate for the memory transaction (block 70). If the CS mode is not interleaved, then the CS range which encompasses the address identifies which chip select signal to activate (block 72).

[0070] It is noted that a similar flowchart may be applicable to channel control circuit 32B, except that decision block 62 checks that the address bit used to select

between the interleaved channels is set. Other embodiments having more than two channels which may be interleaved may use more than one bit, similar to the CS interleaving.

[0071] It is noted that the circuitry performing the flowchart of Fig. 5 may be partially CS generator circuit 50 circuitry and partially transaction queue 30 circuitry. For example, in one embodiment transaction queue 30 issues a memory transaction to the correct channel control circuit 32A-32B based on configuration registers 36A-36B. In such an embodiment, all of the flowchart in Fig. 5 could be implemented in transaction queue 30 and transaction queue 30 may send an indication of the appropriate chip select signal to that channel control circuit 32A-32B. CS generator circuit 50 may be eliminated in such an embodiment. Alternatively, for example, transaction queue 30 could perform blocks 60, 62, 64, and 66 to determine which of channel control circuits 32A-32B to issue the memory transaction to. CS generator circuit 50 may perform blocks 64, 68, 70, and 72. In yet another alternative, transaction queue 30 may broadcast a transaction to both channel control circuits 32A-32B. In such an embodiment, CS generator circuit 50 may perform all of the flowchart in Fig. 5.

[0072] Turning next to Fig. 6, several examples of programmable CS interleaving modes are shown. For each of the examples in Fig. 6, channel interleaving is not selected.

[0073] The first example 80 illustrates a CS interleaving mode of no interleaving. Thus, the channel 0 memory 82 is divided into 4 separate CS regions 84A-84D. The CS regions may be of varying sizes. As illustrated via the arrows pointing to the channel 0 memory 82, the CS_Start for each CS is programmed to the lower end of the range mapped to that CS and the CS_End is programmed to the upper end of that range (and thus may be equal to the CS_Start of the next CS). Also illustrated in example 80 is that the CS regions need not follow the numbering of the CS signals. For example, the CS range corresponding to CS2 is assigned to lower addresses than the CS range corresponding to CS 1. Any order of CS range assignments may be made, as desired.

[0074] The second example 90 illustrates a CS interleaving mode of mixed mode interleaving. Mixed mode CS interleaving refers to the situation in which two or more memory sections corresponding to two or more CS signals are interleaved while other memory sections corresponding to other CS signals are not interleaved. In example 90, CS0 and CS1 are interleaved and CS2 and CS3 are not interleaved. Thus, the channel 0 memory 92 is divided into three CS regions 94A-94C. Like example 82, regions 94A-94B each correspond to a CS (CS3 and CS2, respectively) and the corresponding start and end address fields are programmed as illustrated. On the other hand, region 94C corresponds to both CS 1 and CS0, and thus the start and end address fields for CS 1 and CS0 are both programmed to represent region 94C. Additionally, an address bit is defined

(in configuration register 36AA) to determine the interleaving of the CS0 and CS1 regions. Thus region 94C is divided into subregions (e.g. subregions 96A-96B). The subregions are alternately assigned to CS0 and CS1.

[0075] The third example 100 illustrates a CS interleaving mode of full interleaving, in which all memory sections corresponding to all CS signals are interleaved. Thus, channel 0 memory 102 is a single region. The CS start and end fields are all programmed to represent the addresses of the full memory size. Additionally, two address bits are defined (in configuration register 36AA) to determine the interleaving of the CS0, CS1, CS2, and CS3 regions. Memory 102 is divided into multiple subregions (e.g. subregions 104A-104D). The subregions are alternately assigned to CS0, CS1, CS2, and CS3 as shown.

[0076] Turning now to Fig. 7, an example of channel interleaving is shown. In the example, no CS interleaving is programmed although CS interleaving can be programmed as well, if desired. Additionally, Fig. 7 illustrates an example in which the address bit used to select between the channels is a less significant address bit than the address bits which determine the CS regions. Accordingly, the memory 110 is divided into four separate CS regions 112A-112D. Each CS region is divided into subregions (e.g. subregions 114A and 114B in CS region 112D). The subregions are alternately assigned to channel zero and channel one (Ch0 and Ch1 in Fig. 7).

[0077] Turning next to Fig. 8, several examples of row address, column address, and bank address configurations are shown. More particularly, exemplary values for various configuration registers are illustrated in Fig. 8. The exemplary values are illustrated for an embodiment which uses bit masks to define the portion of the address used for the row address, column address, and bank selection. The examples of Fig. 8 correspond to 40 bit addresses of memory transactions. Other embodiments may have more or fewer address bits, as desired. The least significant bits are on the right in Fig. 8. Additionally, the bit masks are divided (via underscores) into 8 bit portions in Fig. 8 to ease viewing. The underscores are not actually stored in the corresponding configuration registers.

[0078] The first example 120 illustrates a configuration for a 512 Megabyte (MB) CS0 region in which the memory has 4096 rows (4K) and 1024 columns (1K). Thus, the row address comprises 12 bits and the column address comprises 10 bits. Additionally, the CS0 region may include four banks and thus the bank selection includes two bits. The contents of configuration registers 36AE, 36AI, and 36AM are illustrated in first example 120. In first example 120, the bank selection is programmed to select from less significant bits of the address than the column address is selected from, and the column address is selected from less significant bits than the row address is selected from. More particularly, if the least significant bit of the address is referred to as

bit 0, address bits 5 and 6 are used to form the bank selection given the contents of configuration register 36AM in this example. Similarly, address bits 7-16 are used to form the column address given the contents of configuration register 36AI in this example and address bits 17-28 are used to form the row address given the contents of configuration register 36AE in this example.

[0079] A second example 122 is shown having the same CS0 region (512 MB, 4K roses, 1K columns). The contents of configuration registers 36AE, 36AI, and 36AM are illustrated in second example 122 as well. However, in the second example 122, the bank selection is programmed to select from more significant bits than the column address is selected from, but less significant bits than the row address is selected from. More particularly, if the least significant bit of the address is referred to as bit 0, address bits 15 and 16 are used to form the bank selection given the contents of configuration register 36AM in this example. Similarly, address bits 5-14 are used to form the column address given the contents of configuration register 36AI in this example and address bits 17-28 are used to form the row address given the contents of configuration register 36AE in this example.

[0080] A third example 124 is shown having the same CS0 region (512 MB, 4K roses, 1K columns). The contents of configuration registers 36AE, 36AI, and 36AM are illustrated in the third example 124 as well. However, in the third example 124, the column address bit mask includes a gap large enough for the bank selection mask to select bits. Thus, the column address is selected from bits 5 and 8 to 16 (where bit 0 is the least significant bit) of the address of the memory transaction and the bank selection is selected from bits 6 and 7 of the address of the memory transaction. It is noted that, if CS interleaving or channel interleaving were used with CS0, additional bits could be included in the gap in the column address bit mask and the bits which determine which interleaved CS or channel is selected can be bits within the gap as well.

[0081] Examples 120-124 illustrate some of the flexibility available in configuring the memory by configuring the selection of row address, column address, and bank selection. In the configuration of the example 120, the bank selection is selected from the least significant bits out of the row address, column address, and bank selection. Thus, if a set of consecutive memory accesses accessed contiguous cache lines (which are 32 bytes in the illustrated embodiment), the bank selection would change each time before the column address changes. Thus, at least a few of the set of consecutive memory transactions might not encounter an open page, but subsequent transactions of the set might encounter open pages in each bank. On the other hand, the configuration of example 122 may provide an open page for up to 1K contiguous cache lines, all in the same bank. The pages open in other banks may not be affected by the set of consecutive memory transactions. Example

124 would provide an open page for the next consecutive cache line, and then the bank selection would switch. Each example may have advantages for different memory access patterns, and thus may be used based on the expected memory access patterns to CS0.

[0082] A fourth example 126 is shown in Fig. 8 as well. The fourth example is a 1024 MB (1 Gigabyte) CS0 and CS1 interleaved region, where each of the interleaved memory sections is 512 MB with 4K rows and 1K columns. Thus, in example 126 the contents of configuration registers 36AA, 36AE-36AF, 36AI-36AJ, and 36AM-36AN are shown. In the example configuration registers 36AE and 36AF are programmed to the same values, as are configuration registers 36AI and 36AJ and configuration registers 36AM and 36AN. More particularly, the configuration registers are programmed such that the bank is selected from less significant address bits than the column address is selected from, and the column address is selected from less significant address bits than the row address is selected from. However, there is a gap between the address bits selected for the row address and the address bits selected for the column address. The gap is the bit selected, in configuration register 36AA, to be the address bit which selects between CS0 and CS1. Other examples are possible as well, including examples similar to example 122 and example 124. Furthermore, the address bit which selects between CS0 and CS1 could be between the address bits selected for the column address and the address bits selected for the bank selection, or a less significant address bit than the bank selection, etc.

[0083] Turning next to Fig. 9, a block diagram of a carrier medium 300 including a database representative of system 10 is shown. Generally speaking, a carrier medium may include storage media such as magnetic or optical media, e.g., disk or CD-ROM, volatile or non-volatile memory media such as RAM (e.g. SDRAM, RDRAM, SRAM, etc.), ROM, etc., as well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as a network and/or a wireless link.

[0084] Generally, the database of system 10 carried on carrier medium 300 may be a database which can be read by a program and used, directly or indirectly, to fabricate the hardware comprising system 10. For example, the database may be a behavioral-level description or register-transfer level (RTL) description of the hardware functionality in a high level design language (HDL) such as Verilog or VHDL. The description may be read by a synthesis tool which may synthesize the description to produce a netlist comprising a list of gates from a synthesis library. The netlist comprises a set of gates which also represent the functionality of the hardware comprising system 10. The netlist may then be placed and routed to produce a data set describing geometric shapes to be applied to masks. The masks may then be used in various semiconductor fabrication steps to produce a semiconductor circuit or circuits corre-

sponding to system 10. Alternatively, the database on carrier medium 300 may be the netlist (with or without the synthesis library) or the data set, as desired.

[0085] While carrier medium 300 carries a representation of system 10, other embodiments may carry a representation of any portion of system 10, as desired, including any combination of a memory controller, channel control circuit, configuration registers, memory, etc.

[0086] Numerous variations and modifications within the scope of the claimed invention will become apparent to those skilled in the art once the above disclosure is fully appreciated.

Claims

1. A memory controller for coupling to a memory and for coupling to receive an address for access to the memory, the memory controller programmable with an indication of which portion of said address is used to select a storage location in said memory for access in response to said address, and wherein said memory controller is configured to extract said portion of said address for transmission to said memory responsive to said indication.
2. The memory controller as recited in claim 1 comprising one or more registers programmable with the indication.
3. The memory controller as recited in claim 1 or claim 2 wherein said memory is arranged into rows and columns of storage locations, and wherein said portion of said address includes a first portion of said address used to select one of said rows and a second portion of said address used to select one of said columns, and wherein said indication identifies said first portion and said second portion.
4. The memory controller as recited in claim 3 wherein said indication includes a first bit mask, wherein each bit of said first bit mask corresponds to at least one bit of said address, and wherein said circuit is configured to selectively include a first bit of said address in said first portion responsive to a corresponding bit of said first bit mask.
5. The memory controller as recited in claim 4 wherein said indication includes a second bit mask, wherein each bit of said second bit mask corresponds to at least one bit of said address, and wherein said circuit is configured to selectively include a second bit of said address in said second portion responsive to a corresponding bit of said second bit mask.
6. The memory controller as recited in any of claims 3 to 5 wherein said memory is further arranged into banks, and wherein said portion of said address includes a third portion used to select one of said banks, and wherein said indication identifies said third portion.
7. The memory controller as recited in claim 6 wherein said indication comprises a third bit mask, wherein each bit of said third bit mask corresponds to at least one bit of said address, and wherein said circuit is configured to selectively include a third bit of said address in said third portion responsive to a corresponding bit of said third bit mask.
8. The memory controller as recited in any preceding claim wherein said memory is arranged into two or more memory sections, and wherein said memory controller is configured to provide a different select signal to each of said memory sections to select said memory sections to respond to said address, and wherein said memory controller is further programmable with an interleave mode for said memory sections.
9. The memory controller as recited in claim 8 wherein said interleave mode is no interleave.
10. The memory controller as recited in claim 8 wherein said interleave mode is interleave of a subset of said memory sections and no interleave of remaining ones of said memory sections.
11. The memory controller as recited in claim 8 wherein said interleave mode is interleave of said memory sections.
12. The memory controller as recited in any of claims 8 to 11 wherein said portion of said address further includes a fourth portion of said address which is used to select one of said interleaved memory sections, and wherein said indication identifies said fourth portion.
13. The memory controller as recited in claim 12 wherein said indication comprises a fourth bit mask, wherein each bit of said fourth bit mask corresponds to at least one bit of said address, and wherein said circuit is configured to selectively include a fourth bit of said address in said fourth portion responsive to a corresponding bit of said fourth bit mask.
14. The memory controller as recited in any of claims 8 to 13 wherein said memory controller is further programmable with a different page mode policy for each of said memory sections.
15. The memory controller as recited in any preceding claim wherein said memory is arranged in two or more independent memory sections, and wherein said memory controller is configured to provide a

separate channel to each of said memory sections, and wherein said memory controller is further programmable with a channel interleave indication indicating whether or not one or more of said channels are interleaved.

16. The memory controller as recited in claim 15 wherein said channel interleave indication further indicates a fifth portion of said address used to select one of said channels if one or more of said channels are interleaved.

17. A method comprising:

programming a memory controller with an indication of which portion of an address is used to select a storage location in a memory for access in response to said address; receiving said address in said memory controller; and extracting said portion of said address responsive to said indication.

18. The method as recited in claim 17 wherein said memory is arranged into rows and columns of storage locations, and wherein said portion of said address includes a first portion of said address used to select one of said rows and a second portion of said address used to select one of said columns, and wherein said indication identifies said first portion and said second portion.

19. The method as recited in claim 18 wherein said indication includes a first bit mask, wherein each bit of said first bit mask corresponds to at least one bit of said address, the method further comprising selectively including a first bit of said address in said first portion responsive to a corresponding bit of said first bit mask.

20. The method as recited in claim 19 wherein said indication includes a second bit mask, wherein each bit of said second bit mask corresponds to at least one bit of said address, the method further comprising selectively including a second bit of said address in said second portion responsive to a corresponding bit of said second bit mask.

21. The method as recited in any of claims 18 to 20 wherein said memory is further arranged into banks, and wherein said portion of said address includes a third portion used to select one of said banks, and wherein said indication identifies said third portion, the method further comprising transmitting said third portion separate from said first portion and said second portion.

22. The method as recited in claim 21 wherein said

indication comprises a third bit mask, wherein each bit of said third bit mask corresponds to at least one bit of said address, the method further comprising selectively including a third bit of said address in said third portion responsive to a corresponding bit of said third bit mask.

23. The method as recited in any of claims 17 to 22 wherein said memory is arranged into two or more memory sections, the method further comprising:

providing a different select signal to each of said memory sections to select said memory sections to respond to said address; and programming said memory controller with an interleave mode for said memory sections.

24. The method as recited in claim 23 wherein said interleave mode is no interleave.

25. The method as recited in claim 23 wherein said interleave mode is interleave of a subset of said memory sections and no interleave of remaining ones of said memory section.

26. The method as recited in claim 23 wherein said interleave mode is interleave of said memory sections.

27. The method as recited in any of claims 23 to 26 wherein said portion of said address further includes a fourth portion of said address which is used to select one of said interleaved memory sections, and wherein said indication identifies said fourth portion.

28. The method as recited in claim 27 wherein said indication comprises a fourth bit mask, wherein each bit of said fourth bit mask corresponds to at least one bit of said address, the method further comprising selectively including a fourth bit of said address in said fourth portion responsive to a corresponding bit of said fourth bit mask.

29. The method as recited in any of claims 23 to 28 further comprising programming said memory controller with a different page mode policy for each of said memory sections.

30. The method as recited in any of claims 17 to 29 wherein said memory is arranged in two or more independent memory sections, the method further comprising:

providing a separate channel to each of said memory sections; and programming said memory controller with a channel interleave indication indicating wheth-

er or not one or more of said channels are interleaved.

31. The method as recited in claim 31 wherein said channel interleave indication further indicates a fifth portion of said address used to select one of said channels if one or more of said channels are interleaved. 5
32. A circuit defining mechanism comprising one or more databases representing the memory controller as recited in any of claims 1-16. 10
33. The circuit defining mechanism as recited in claim 32 wherein the one or more databases include one or more netlists representing the memory controller. 15
34. A carrier medium carrying the circuit defining mechanism as recited in claim 32. 20

25

30

35

40

45

50

55




Fig. 1




Fig. 2




Fig. 3

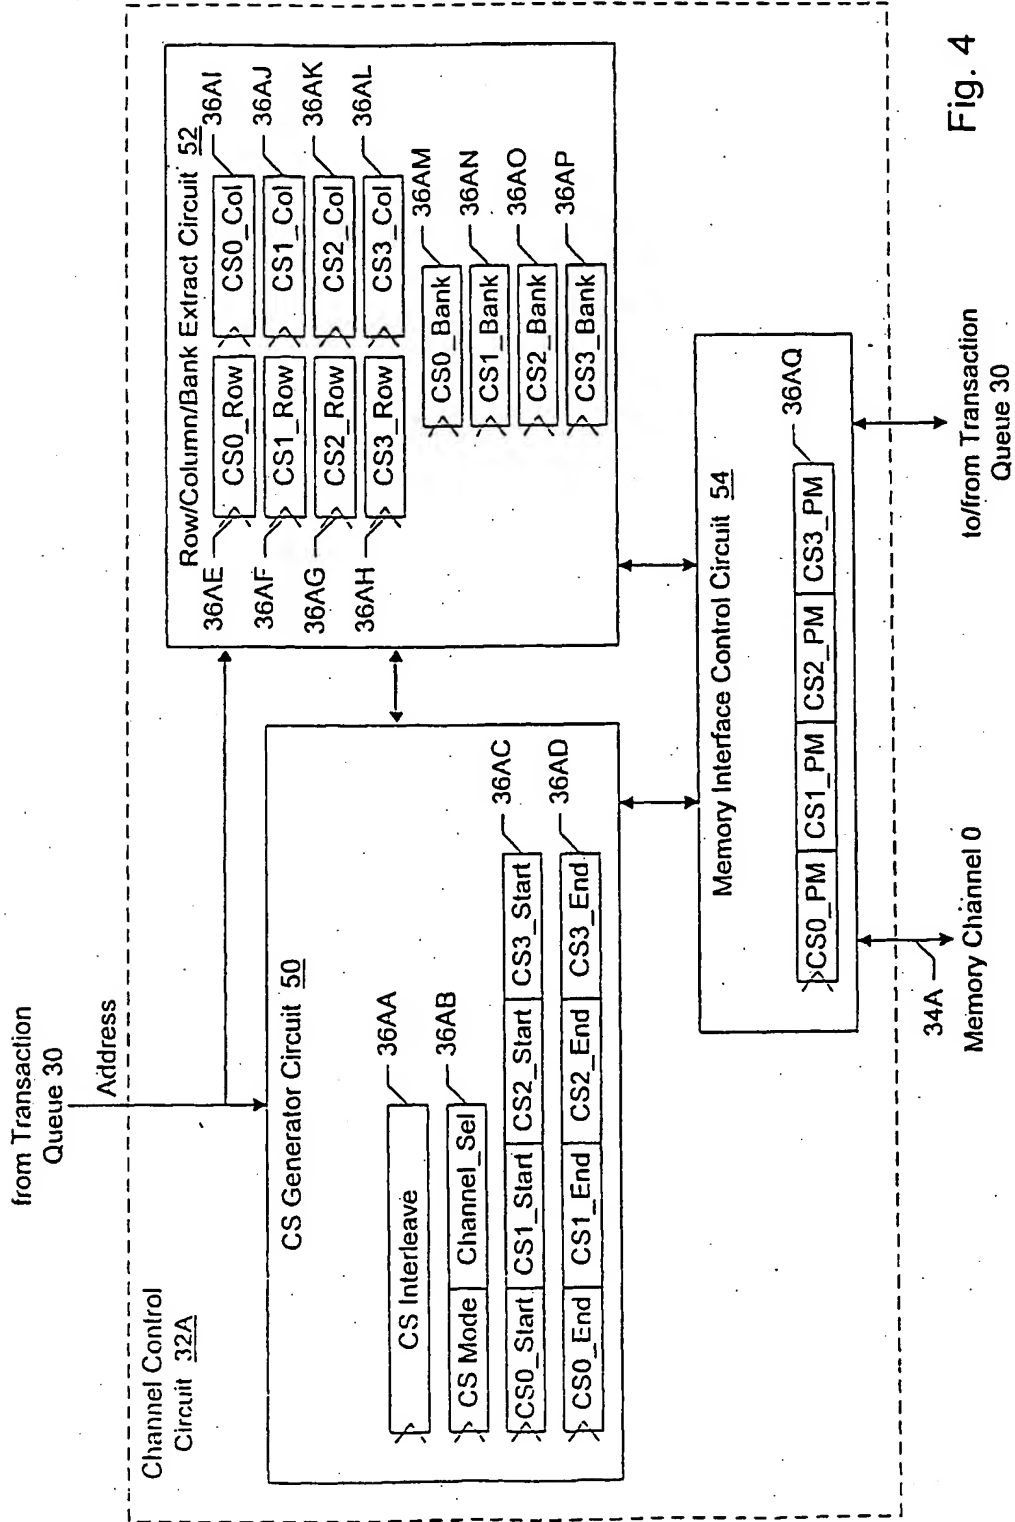


Fig. 4

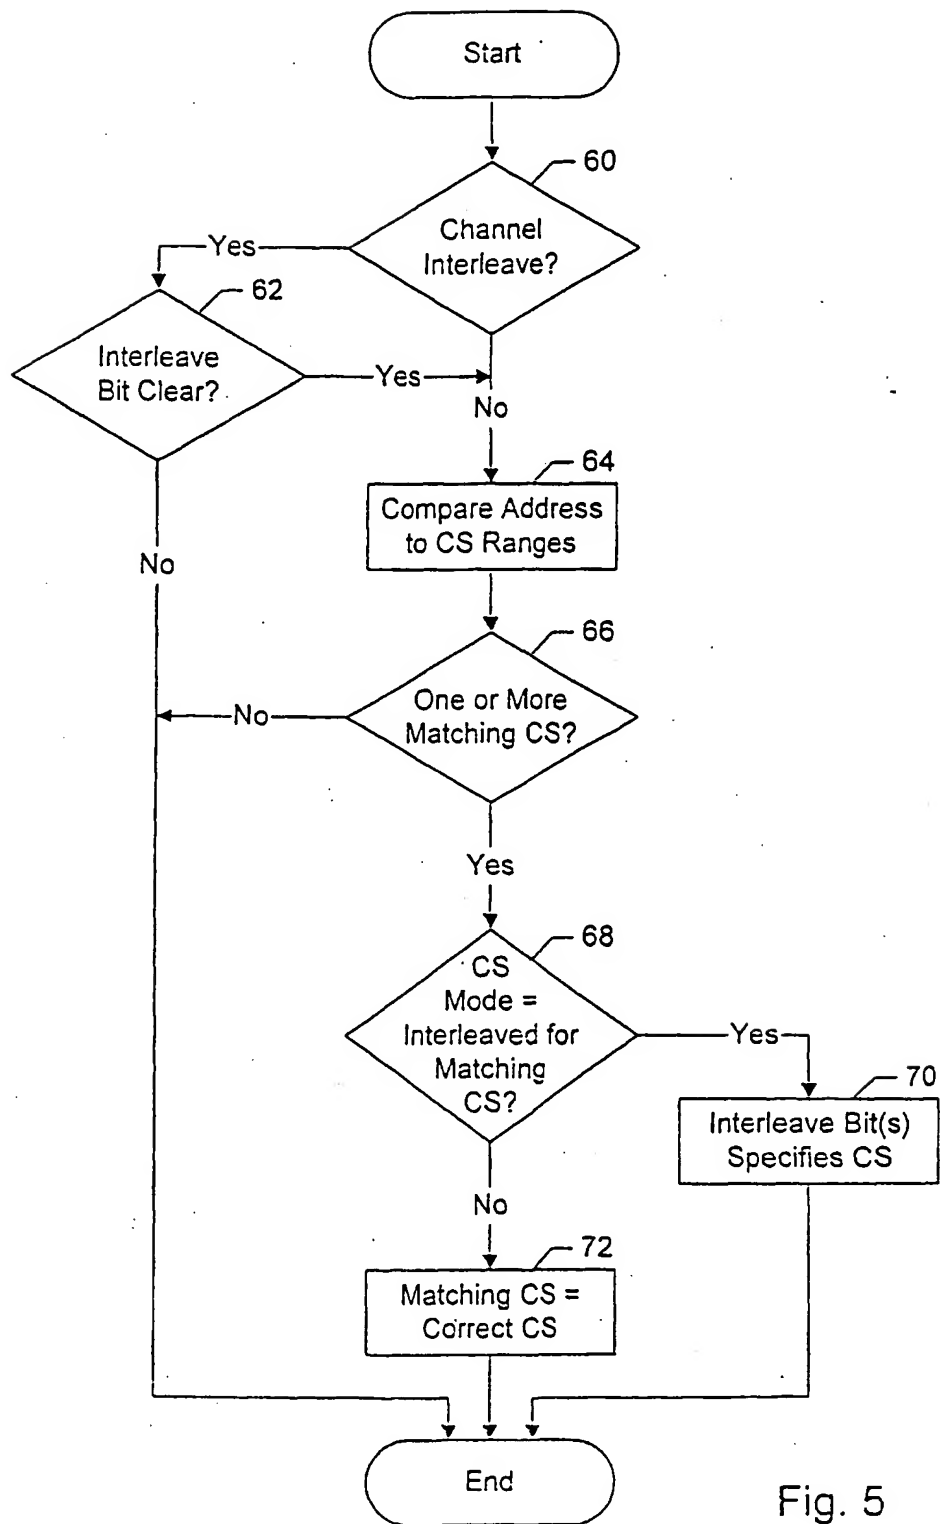


Fig. 5

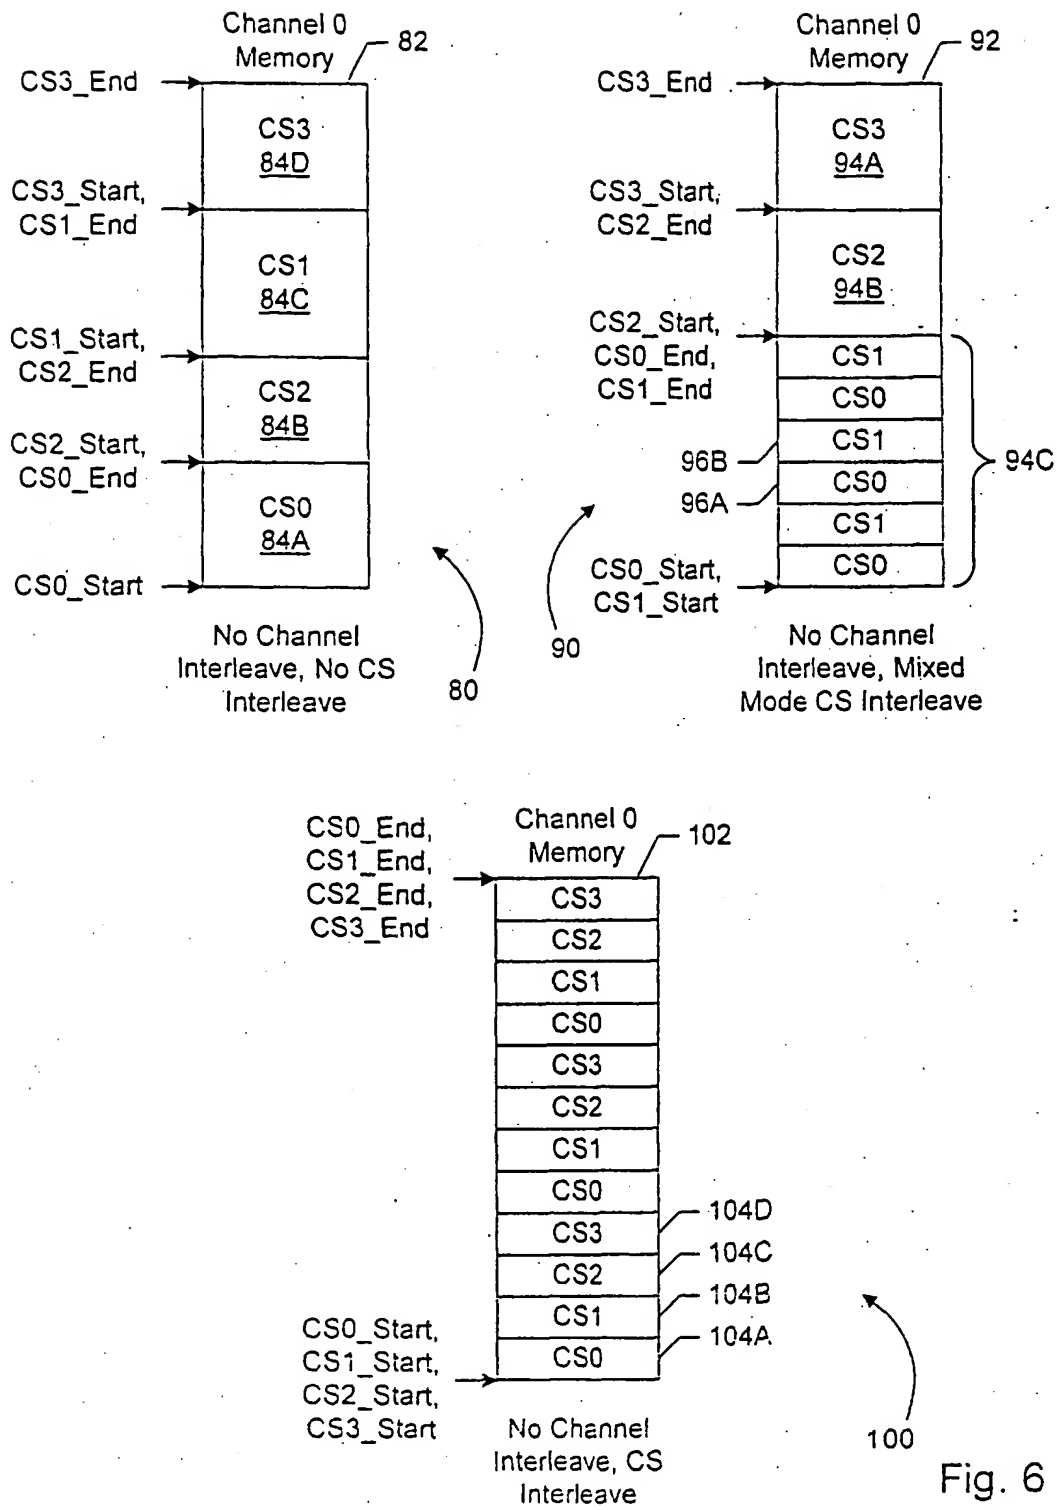


Fig. 6

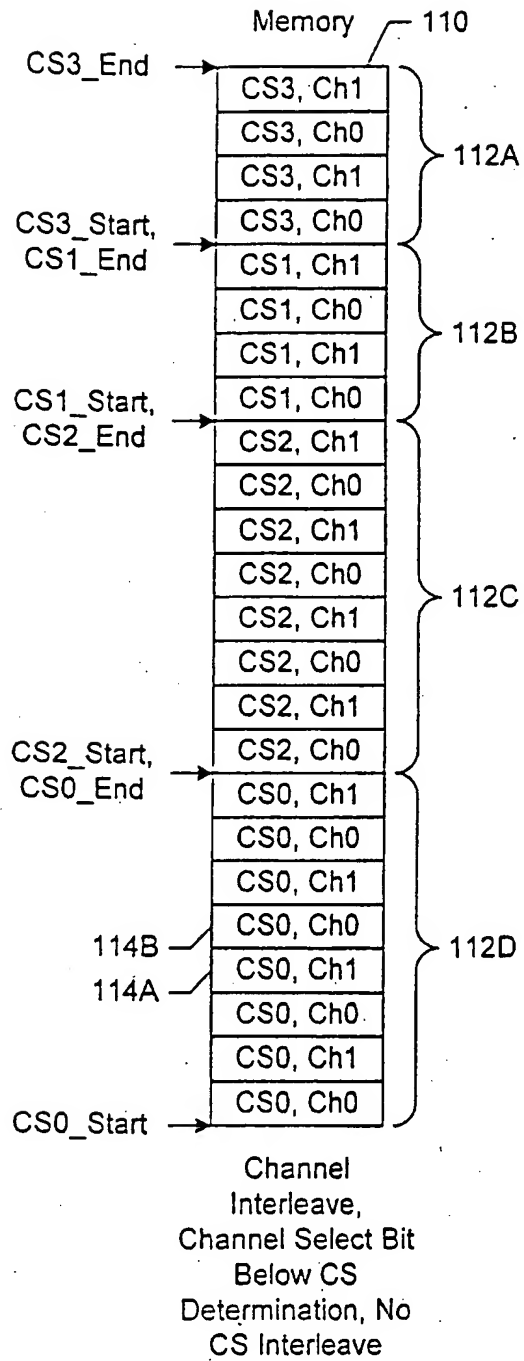


Fig. 7

512 MB CS0 Region, 4K Rows, 1K Columns

| | |
|----------------------|--|
| CS0_Row <u>36AE</u> | 00000000_00011111_11111110_00000000_00000000 |
| CS0_Col <u>36AI</u> | 00000000_00000000_00000001_11111111_10000000 |
| CS0_Bank <u>36AM</u> | 00000000_00000000_00000000_00000000_01100000 |

120

512 MB CS0 Region, 4K Rows, 1K Columns, Alternative Encoding

| | |
|----------------------|--|
| CS0_Row <u>36AE</u> | 00000000_00011111_11111110_00000000_00000000 |
| CS0_Col <u>36AI</u> | 00000000_00000000_00000000_01111111_11100000 |
| CS0_Bank <u>36AM</u> | 00000000_00000000_00000001_10000000_00000000 |

122

512 MB CS0 Region, 4K Rows, 1K Columns, 64 Byte Bank Interleave

| | |
|----------------------|--|
| CS0_Row <u>36AE</u> | 00000000_00011111_11111110_00000000_00000000 |
| CS0_Col <u>36AI</u> | 00000000_00000000_00000001_11111111_00100000 |
| CS0_Bank <u>36AM</u> | 00000000_00000000_00000000_00000000_11000000 |

124

1024 MB CS0 and CS1 Region, 4K Rows, 1K Columns

| | |
|---------------------------|--|
| CS_Interleave <u>36AA</u> | 00000000_00000000_00000010_00000000_00000000 |
| CS0_Row <u>36AE</u> | 00000000_00111111_11111100_00000000_00000000 |
| CS0_Col <u>36AI</u> | 00000000_00000000_00000001_11111111_10000000 |
| CS0_Bank <u>36AM</u> | 00000000_00000000_00000000_00000000_01100000 |
| CS1_Row <u>36AF</u> | 00000000_00111111_11111100_00000000_00000000 |
| CS1_Col <u>36AJ</u> | 00000000_00000000_00000001_11111111_10000000 |
| CS1_Bank <u>36AN</u> | 00000000_00000000_00000000_00000000_01100000 |

126

Fig. 8




Fig. 9

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:
12.05.2004 Bulletin 2004/20

(51) Int Cl.7: **G06F 12/06**

(43) Date of publication A2:
27.03.2002 Bulletin 2002/13

(21) Application number: **01308014.8**

(22) Date of filing: **20.09.2001**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR**
Designated Extension States:
AL LT LV MK RO SI

• **Hayter, Mark D.**
Menlo Park, California 94025 (US)
• **Cho, James Y.**
Los Gatos, California 95032 (US)

(30) Priority: **20.09.2000 US 665989**

(71) Applicant: **Broadcom Corporation**
Irvine, California 92618 (US)

(74) Representative:
Jehle, Volker Armin, Dipl.-Ing. et al
Patentanwälte
Bosch, Graf von Stosch, Jehle,
Flüggensstrasse 13
80639 München (DE)

(72) Inventors:
• **Keller, James B.**
Palo Alto, California 94303 (US)

(54) **Memory controller with programmable configuration**

(57) A memory controller provides programmable flexibility, via one or more configuration registers, for the configuration of the memory. The memory may be optimized for a given application by programming the configuration registers. For example, in one embodiment, the portion of the address of a memory transaction used to select a storage location for access in response to the memory transaction may be programmable. In an implementation designed for DRAM, a first portion may be programmably selected to form the row address and a second portion may be programmably selected to form

the column address. Additional embodiments may further include programmable selection of the portion of the address used to select a bank. Still further, interleave modes among memory sections assigned to different chip selects and among two or more channels to memory may be programmable, in some implementations. Furthermore, the portion of the address used to select between interleaved memory sections or interleaved channels may be programmable. One particular implementation may include all of the above programmable features, which may provide a high degree of flexibility in optimizing the memory system.




Fig. 1



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 8014

| DOCUMENTS CONSIDERED TO BE RELEVANT | | | |
|---|--|--|--|
| Category | Citation of document with indication, where appropriate, of relevant passages | Relevant to claim | CLASSIFICATION OF THE APPLICATION (Int.Cl.7) |
| X | US 5 412 788 A (COLLINS HANSEL A ET AL) 2 May 1995 (1995-05-02) | 1,2,17 | G06F12/06 |
| Y | * figure 4C * * column 11, line 38 - line 44 * * column 12, line 30 - line 59 * * column 13, line 59 - line 64 * * column 14, line 58 - line 64 * | 3-16, 18-34 | |
| Y | GB 2 256 293 A (RESEARCH MACHINES PLC) 2 December 1992 (1992-12-02) * page 1, line 26 - line 27 * * page 2, line 1 - line 19 * * page 3, line 19 - line 27 * * page 4, line 22 - line 24 * * figures 4-6 * | 3-16, 18-34 | |
| | | | TECHNICAL FIELDS SEARCHED (Int.Cl.7) |
| | | | G06F |
| The present search report has been drawn up for all claims | | | |
| Place of search THE HAGUE | | Date of completion of the search 9 March 2004 | Examiner Jardon, S |
| <p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p> | | | |

EPO FORM 1500 (03.02) (P04C01)

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 01 30 8014

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

09-03-2004

| Patent document cited in search report | | Publication date | Patent family member(s) | Publication date |
|---|---|---------------------|----------------------------|---------------------|
| US 5412788 | A | 02-05-1995 | NONE | |
| GB 2256293 | A | 02-12-1992 | NONE | |

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82